Learning to Suppress Competing Information:
Do the Skills Transfer?

Morton Ann Gernsbacher
University of Wisconsin-Madison

Research and Advanced Concepts Office

October 2001

U.S. Army Research Institute
for the Behavioral and Social Sciences

20011102 060

**U.S. Army Research Institute
for the Behavioral and Social Sciences**

A Directorate of the U.S. Total Army Personnel Command

EDGAR M. JOHNSON
Director

NOTICES

REPORT DOCUMENTATION PAGE

1. REPORT DATE (dd-mm-yy) October 2001	2. REPORT TYPE Final	3. DATES COVERED (from. . . to) August 8, 1998—December 31, 1999	
4. TITLE AND SUBTITLE Learning to Suppress Competing Information: Do the Skills Transfer?		5a. CONTRACT OR GRANT NUMBER DASW01-98-M-2299	
		5b. PROGRAM ELEMENT NUMBER 2O161102B	
6. AUTHOR(S) Morton Ann Gernsbacher		5c. PROJECT NUMBER B74F	
		5d. TASK NUMBER 1901	
		5e. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of Psychology University of Wisconsin-Madison Madison WI 53706		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Institute for the Behavioral and Social Sciences 5001 Eisenhower Avenue Alexandria, VA 22333-5600		10. MONITOR ACRONYM ARI	
		11. MONITOR REPORT NUMBER Research Note 2002-01	

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT *(Maximum 200 words)*:

This report discusses the laboratory research on the cognitive mechanism of suppression. The goal of this study is to identify the cognitive processes and mechanisms that underlie language comprehension and comprehension in general.

15. SUBJECT TERMS
Cognitive processes, language comprehension, training, skills transfer, suppression, lexical access, anaphoric reference and Cataphoric Reference.

SECURITY CLASSIFICATION OF:			19. LIMITATION OF ABSTRACT	20. NUMBER OF PAGES	21. RESPONSIBLE PERSON (Name and Telephone Number)
16. REPORT Unclassified	17. ABSTRACT Unclassified	18. THIS PAGE Unclassified	Unlimited	33	Michael Drillings, Ph.D. (703)617-8641

i

Learning to Suppress Competing Information: Do the Skills Transfer?

Abstract

This proposal sought funds to continue to conduct laboratory research on the cognitive mechanism of suppression. Suppression is conceptualized as an active dampening of the automatic activation of mental representations. Thus, suppression attenuates the interference caused by the activation of inappropriate or irrelevant information. In my previous research (much of which was supported by previous ARI funds), I have empirically illustrated the crucial role that suppression plays in many cognitive tasks. Furthermore, during a previous funding period I discovered that (1) the mechanism of suppression is under strategic control, and (2) persons can be taught to suppress competing information. The next stage of research greatly extended these two recent discoveries by answering the following question: *Does training in suppression of one type of competing information transfer to skill in suppressing another type of competing information?* This question was answered through a series of laboratory experiments. The results of these experiments inform us about the transferability of training of the crucial skill of suppression. Discovering that training in suppression in one domain leads to improved suppression in another domain has great theoretical and practical implications. At the theoretical level, discovering that training in suppression in one domain leads to improved suppression in another domain supports the hypothesis of *one* general, cognitive mechanism of suppression. At the practical level, discovering that training in suppression in one domain leads to improved suppression in another domain demonstrates that persons' ability to suppress information can be improved, even without specific training in the domain in which suppression is required.

Morton Ann Gernsbacher
Department of Psychology
University of Wisconsin-Madison
Madison WI 53706

Learning to Suppress Competing Information: Do the Skills Transfer?

The goal of my research is to identify the cognitive processes and mechanisms that underlie language comprehension and comprehension in general. I have identified a few of those processes and mechanisms in a framework I call the Structure Building Framework (Gernsbacher, 1990; 1991; 1995; 1997). According to the Structure Building Framework, the goal of comprehension is to build coherent mental representations or *structures*. These structures represent clauses, sentences, paragraphs, passages, and other meaningful units.

To build these structures, first, comprehenders lay foundations for their mental structures (Carreiras, Gernsbacher, & Villa, 1995; Gernsbacher & Hargreaves, 1988; Gernsbacher & Hargreaves, 1992; Gernsbacher, Hargreaves, & Beeman, 1989). Then comprehenders develop their mental structures by mapping on information, when that incoming information coheres or relates to the previous information (Deaton & Gernsbacher, in press; Gernsbacher, 1996; Gernsbacher & Givón, 1995; Gernsbacher, Hallada, & Robertson, in press; Gernsbacher & Robertson, 1992; Haenggi, Gernsbacher, & Bolliger, 1993; Haenggi, Kintsch, & Gernsbacher, 1995; Traxler & Gernsbacher, 1995). However, if the incoming information is less coherent, comprehenders employ a different process: They shift and initiate a new substructure (Foertsch & Gernsbacher, 1994; Gernsbacher, 1985). So, most mental representations comprise several branching substructures.

The building blocks of mental structures are what I refer to as memory nodes. According to the Structure Building Framework, memory nodes are activated by incoming stimuli. Once activated, the information they represent can be used by cognitive processes. Furthermore, according to the Structure Building Framework, activated memory nodes transmit processing signals. These processing signals either suppress or enhance the activation of other memory nodes. In other words, once memory nodes are activated, two mechanisms modulate their level of activation. The two mechanisms are suppression and enhancement.

Suppression decreases or dampens the activation of memory nodes when the information they represent is no longer as necessary for the structure being built. Enhancement increases or boosts the activation of memory nodes when the information they represent is relevant to the structure being built. By modulating the activation of memory nodes, suppression and enhancement contribute greatly to language comprehension. I want to stress, however, that suppression and enhancement are general cognitive mechanisms. They are not dedicated solely to language; they play vital roles in nonlinguistic processes, too. But language comprehension draws heavily on these two mechanisms. The proposed research focuses on the mechanism of suppression.

While I believe that most people can appreciate that we need a mechanism that enhances relevant or related information, I have suggested that a mechanism that suppresses inappropriate or irrelevant information is perhaps even more crucial to the goal of comprehension, or in the words of the Structure Building Framework, the goal of building coherent mental representations. The reason why we need a mechanism of suppression is that whenever we comprehend language, we experience various types of interference. Sometimes this interference during comprehension arises from the external environment, as when we conduct a conversation in a noisy restaurant or listen to a lecture with some clod in the audience whispering beside us.

Other times interference during comprehension arises internally, as when we have to deal with the competing meanings of a word or phrase, or the alternate references of a pronoun. Indeed, even in a process as seemingly straightforward as reading a string of letters, such as ROWS, mental information that is related to that string of letters is often activated in our minds.

This mental information might be orthographically related (such as the letter string BOWS), or phonologically related (such as the sound /roz/), or even semantically related (such as the concept "rose"). And indeed, laboratory experiments demonstrate that adults have difficulty quickly rejecting the letter string ROWS as not being a member of the semantic category, *flower* (van Orden, 1987; van Orden, Johnston, & Hale, 1988).

External information often interferes with our comprehension. For example, laboratory experiments demonstrate that it is harder to read a word when it is written within a line-drawing of an object, and it is harder to name a line-drawn object if a word is written within it (Smith & McGee, 1980; Rayner & Posnansky, 1978; Rosinski, Golinkoff, & Kukish, 1975). Thus, successful comprehension involves successfully attenuating or inhibiting interfering information. This attenuation of interfering information during comprehension is the type of inhibition that I have investigated in my research and propose to continue to investigate.

I have argued that a particular cognitive mechanism, what I call the cognitive mechanism of suppression, reduces such interference. In my previous research (much of which was supported by previous ARI funds), I have empirically illustrated the crucial role that suppression plays in many comprehension phenomena. These phenomena include the following (the **boldfaced** citations are work supported by previous ARI funds):

LEXICAL ACCESS — how comprehenders understand or "access" from their memory the meanings of words (**Faust, Balota, Duchek, Gernsbacher & Smith, 1997**; **Faust & Gernsbacher, 1996**; **Gernsbacher & Faust, 1991b**; **Gernsbacher & Faust, 1995**; **Gernsbacher & St. John, in press**);

ANAPHORIC REFERENCE — how comprehenders understand to whom or what anaphors, like pronouns refer (**Foertsch & Gernsbacher, 1994**; Gernsbacher, 1989; **Garnham, Traxler, Oakhill, & Gernsbacher, 1996**);

CATAPHORIC REFERENCE — how words that are marked by devices, such as spoken stress, gain a privileged status in comprehenders' memories (**Gernsbacher & Jescheniak, 1995**; Gernsbacher & Shroyer, 1989);

SYNTACTIC PARSING — how we decode the grammatical forms of sentences into meaning (**Gernsbacher & Robertson, 1997**; **Gernsbacher & Shlesinger, in press**);

SURFACE INFORMATION LOSS — the finding that seemingly superficial information, such as syntactic form, is often forgotten more rapidly than seemingly more important information, such as thematic content (Gernsbacher, 1985; **Gernsbacher & Shlesinger, in press**);

METAPHOR INTERPRETATION — how we understand figurative expressions such as "lawyers are sharks" (**Gernsbacher, Keysar, & Robertson, 1995**; **Gernsbacher & Robertson, in press**);

INFERENCING — how comprehenders infer information that is only implied by a text or discourse (Carreiras & Gernsbacher, 1992; Gernsbacher, 1991b; **Gernsbacher, 1994**; Gernsbacher & Robertson, 1992; Gernsbacher, Goldsmith, & Robertson, 1992; Oakhill, Garnham, Gernsbacher, & Cain, 1992); and

GENERAL COMPREHENSION SKILL — skill at comprehending linguistic and nonlinguistic media (Gernsbacher, 1993; **Gernsbacher, in press-a**; **Gernsbacher, in press-b**; Gernsbacher & Faust, 1991a; **Gernsbacher & Faust, 1995**; **Gernsbacher & Robertson, 1995**; Gernsbacher, Varner, & Faust, 1990).

In the next section, I shall briefly review some of these experiments. They demonstrate that the mechanism of suppression plays a powerful role in many language comprehension phenomena. Indeed, the role is so crucial that persons who are less skilled at comprehension are marked by less efficiency in suppressing or inhibiting interfering information. Let me begin by describing the role of suppression in lexical access.

The Role of Suppression in Lexical Access

During lexical access, the cognitive mechanism of suppression attenuates the interference caused by other lexical information that is activated when a printed word is read, or a spoken word is heard. This information might be the meanings of a word that are not relevant to the immediate context — for example, the saloon meaning of *bar* in the pun *Two men walk into a bar and a third man ducks.* Or the interfering information might be other words or phrases that are related to the sound pattern of a spoken word or phrase, as in the classic *new display* often erroneously interpreted as *nudist play.*

Most models of lexical access propose that multiple types of information are activated when we read or hear a word; however, my research demonstrates that the mechanism of suppression dampens the activation of the unnecessary information. To examine the role of suppression in lexical access, I have capitalized on a phenomenon that I believe is a quintessential demonstration of the activation of superfluous information during lexical access. The phenomenon involves the comprehension of homonyms — words that share the same lexical form but differ in their meaning or origin, for example, *chest, tire, bowl, match, organ, head, plot, ring, nail.* All languages have homonyms (e.g., *cola, sal, calle, bota, trompa* in Spanish). Indeed, they are usually the most frequently occurring words in a language. Dahlgren has suggested that the average English word has 3 different meanings, and Britton has estimated that homonyms comprise about 40% of our most common open class words. Indeed, the more common the word, the more likely it is to have multiple meanings.

The homonym phenomenon that I have empirically studied is this: Immediately after we hear or read a homonym such as *match* or *duck*, multiple meanings are activated. And more intriguingly, this activation of multiple meanings occurs, regardless of the semantic or syntactic context in which the homonym occurs. For example, immediately after we hear or read, the homonym, *match*, in the sentence, *He lit the match*, both the "firestick" and the "competition" meaning are activated. Immediately after we hear or read the homonym, *duck*, in the sentence, *He needed to duck*, both the "crouching" and the University of Oregon mascot meaning are activated.

Swinney (1979) and Tanenhaus, Leiman, and Seidenberg (1979) were among the first researchers to demonstrate this non-intuitive phenomenon, more than 15 years ago; it has been replicated numerous times since. These researchers also demonstrated — in line with our introspections — that the contextually inappropriate meanings of homonyms do not remain activated forever. What happens to these contextually inappropriate meanings? How do they become less activated? I have proposed that the cognitive mechanism of suppression dampens

their activation. More specifically I have hypothesized that memory nodes that represent a higher-level structure — in this case the sentence-level structure — transmit processing signals to suppress the activation of the inappropriate lexical-level meanings. My research has provided several sources of converging evidence to support this proposal.

For example, in a series of laboratory experiments, postdoctoral fellow Mark Faust and I empirically demonstrated that suppression and not decay reduces the activation of inappropriate meanings (Gernsbacher & Faust, 1991b). That is, inappropriate meanings do not lose activation over time simply because their activation fades with time. We also empirically ruled out a mental "winner takes all" explanation: When inappropriate meanings become less activated it is not because the more appropriate meanings have become more activated. In other words, the inappropriate and appropriate meanings are not "slugging it out mano a mano"; rather the source of the activation reduction comes from a higher-level.

Indeed, using a parallel distributed processing network, postdoctoral fellow Mark St. John and I computationally demonstrated how sentence-level suppression can dampen the activation of contextually inappropriate word meanings (Gernsbacher & St. John, in press). In our connectionist network, suppression driven by a sentence-level representation, what St. John refers to as a gestalt level of representation, was the only type of top-down feedback we allowed, and that alone allowed us to perfectly simulate the behavioral data.

Further demonstrating that suppression and not simply decay is the mechanism responsible for decreasing the activation of the inappropriate meanings of homonyms, laboratory coordinator Rachel Robertson and I empirically demonstrated that suppression carries costs (Gernsbacher & Robertson, 1994). After subjects read a sentence such as *He lit the match* they were considerably slower and considerably less accurate at simply verifying that the sentence *He won the match* made sense. If after reading the *He lit the match* sentence, the inappropriate meaning of match simply decayed, that is, the competition meaning of match simply returned to baseline, that meaning should not have been harder to activate in order to comprehend the subsequent sentence.

Furthermore, as I shall describe later in this section, I have conducted an extensive series of experiments demonstrating that individuals who are less efficient at suppressing many types of information, for example, the color of ink in a Stroop color naming task, hold onto inappropriate meanings considerably longer than do individuals who are more efficient in suppressing extraneous information. And most recently postdoctoral fellow Faust and I discovered a right-visual field, left-cerebral hemisphere advantage for suppressing the inappropriate meanings of homonyms (Faust & Gernsbacher, 1996). When we presented homonyms to the left-visual field (thereby hypothetically stimulating the right-hemisphere prior to the left-hemisphere), resolution of homonym meanings was slightly delayed. Although we still have miles to go before being able to stake our explorers' flag atop the cerebral location of our putative suppression mechanism, we find it less plausible that a decay mechanism would be similarly lateralized.

From all of these findings, I conclude that the mechanism of suppression, which enables the attenuation of interfering mental activation, such as the inappropriate meanings of homonyms, plays a crucial role in lexical access. I shall turn now to review the research I have conducted that investigates the role of suppression in anaphoric reference.

The Role of Suppression in Anaphoric Reference

Anaphoric reference is the process by which readers or listeners understand to whom or to what an anaphor, such as a pronoun, refers. In a series of experiments, I discovered that suppression enables anaphoric reference by attenuating the interference caused by the activation

of other referents. By other referents I mean the people or things to whom or which an anaphoric expression does not refer. For example, consider the sentence, *Ann predicted that Pam would lose the track race, but she came in first very easily.* In this sentence, the pronoun *she* is an anaphoric device, which most people interpret to refer to the referent *Pam.* I discovered that correctly interpreting such anaphoric devices is not so much a matter of activating one of the two possible referents: Both are highly activated because they were just mentioned in the first clause. Rather, understanding to whom the pronoun *she* in the second clause refers, depends on how quickly comprehenders can reduce the activation of the referent to whom the pronoun *she* does not refer (i.e., *Ann* in the example sentence).

In my experiments subjects read sentences word by word. The first clause of each sentence introduced two participants, for example, *Ann* and *Pam* as in *Ann predicted that Pam would lose the track race.* In the second clause, one of those two participants was referred to anaphorically, using either a very explicit repeated name anaphor, such as *Pam*, or a less explicit pronominal anaphor, such as *she,* as in *but she came in first very easily.* I measured activation of the anaphors' referents (like *Pam)* and what I called the nonreferents (like *Ann)* using the probe verification task. Subjects were shown a test name, like Pam or Ann, or a name that didn't occur in the sentence, and their task was to verify whether the test name had occurred in the sentence. Presumably, the faster subjects respond to the test name, the more activated the participant represented by that test name is. In half the experimental sentences the referent was the first-mentioned participant, and in half the referent was the second-mentioned participant, as *Pam* was in the example sentence. In my first experiment I measured activation immediately before versus immediately after the name versus pronoun anaphors occurred. The first test point served as a baseline. Figure 1 (on the next page) displays the data.

As Figure 1 illustrates, immediately after the very explicit name anaphors were read, the referents were considerably more activated than they were before: that is, reaction times decreased. More intriguingly, immediately after the very explicit name anaphors were read, the nonreferents were considerably less activated than they were before; that is, reaction times increased. However, as Figure 1 illustrates, this pattern occurred only for the very explicit name anaphors. For the pronouns, neither the referents nor the nonreferents changed in the activation (indeed, the data were so similar that the two lines lie on top of one another). These data suggest that very explicit repeated name anaphors immediately lead to both enhancement of their referents and suppression of nonreferents. In contrast, less explicit — and indeed momentarily ambiguous — pronoun anaphors do not immediately lead to either suppression or enhancement.

Figure 1

This pattern has been replicated in English, Spanish, and Korean. Furthermore, this pattern demonstrates a relation between anaphoric markedness and the mechanisms of suppression and enhancement. The more marked an anaphor is, the more its referent will be enhanced and the more other concepts (nonreferents) will be suppressed. Why are the mechanisms of suppression and enhancement a function of anaphoric markedness? According to the Structure Building Framework, suppression and enhancement signals are transmitted by activated memory nodes. So, with anaphoric reference, suppression and enhancement signals could be triggered by activated memory nodes that represent information about the referent's identity. And the most available source of information about a referent's identity comes from the anaphor itself.

Repeated proper names, which are more marked, provide quite a bit of information about their referents. Indeed, repeated proper names usually have only one clear referent. We rarely have to say *The Quentin A. Summerfield, who lives on Adams Street.* In between the extremes of repeated proper names and relatively ambiguous pronouns is another type of anaphor: common noun anaphors, such as *the bird* in *John fed the robin. The bird was hungry.* Common noun anaphors are often superordinate categories of their referents. The anaphor *the bird* could refer to a robin, a sparrow, an ostrich, or an emu. Thus, common noun anaphors provide less information about their referents than repeated names do, but more information than pronouns do. Common noun anaphors, should therefore, lead to a medium amount of suppression and enhancement, which is exactly what laboratory experiments demonstrate.

Figure 2 displays the amount of suppression and enhancement that I observed immediately following repeated name anaphors and pronouns. The amount of suppression is indicated by filled bars, and the amount of enhancement is indicated by unfilled bars. Figure 2 also displays the amount of suppression and enhancement observed in an experiment by Dell, McKoon, and Ratcliff (1983), who examined common noun anaphors. Notice that the amount of suppression and enhancement triggered by the common noun anaphors was less than that triggered by the repeated names. And the amount of suppression and enhancement triggered by the gender-ambiguous pronouns that I investigated in the experiment I just described was even less. Thus, the more marked the anaphor, the more likely it is to trigger the suppression of other concepts and the enhancement of its own referent.

Information about a referent's identity also comes from sources outside the anaphor, for instance, semantic, syntactic and pragmatic context. I predicted that the memory nodes representing these other sources of information also trigger suppression and enhancement signals, but they do so more slowly and less powerfully. I tested this prediction in three further experiments. In one experiment, I measured activation immediately before repeated-name versus pronoun anaphors, as I did before, and again this before-the-anaphor test point served as a baseline. However, in this experiment my comparison test point was at the END of the sentence, AFTER the semantic/pragmatic information, which could disambiguate the syntactically ambiguous pronouns, had occurred. For example, activation was measured at the two test points indicated by asterisks in the following example sentence: *Ann predicted that Pam would lose the track race, but * Pam / she came in first very easily. * By the end of the sentence, even the gender-ambiguous pronoun anaphors had led to a reliable amount of suppression of the nonreferents.

In a further experiment, I placed the contextual information before the anaphors, as in, *Ann lost a track race to Pam. Enjoying the victory,* or *Accepting the defeat, Pam/Ann/she headed toward the shower.* Despite the context preceded the anaphors, the less-explicit pronoun anaphors still did not lead to a reliable amount of suppression until the end of the sentence. Thus, information from outside an anaphor can also trigger suppression and enhancement, although it does so more slowly and less powerfully. This is good, because with zero anaphors, as in *Ann lost a tennis match to Pam and 0 cried all the way home,* the anaphor provides no information about its referent. All the information is provided by the semantic, pragmatic, and syntactic context. Therefore, zero anaphors should be the least effective at triggering suppression. This prediction was confirmed in an experiment by Corbett and Chang (1983).

Figure 2

More Explicit

Repeated Names
Common Nouns
Ambiguous Pronouns

0   20   40   60   80   100   120   140

Less Explicit   ms

☐ *Enhancement* ■ *Suppression*

Figure 3

More Explicit

Repeated Nouns
Pronouns
Zero Anaphors

900   950   1000   1050   1100

**Less Explicit**   Reaction Time (ms)

Figure 3 displays an estimate of the amount of suppression triggered by three types of anaphors, repeated name anaphors, pronouns that were disambiguated by the end of the sentence, and zero anaphors. This estimate is based on Corbett and Chang's (1983) subjects' verification times to nonreferents. The slower responses to the nonreferents are, the more they have been suppressed. (All of these measurements were taken at the end of a sentence, after the disambiguating information had occurred.) As Figure 3 illustrates, these data again illustrate a markedness relation: The more marked the anaphors, from repeated names, to semantically-cued pronouns, to zero anaphors, the more likely they are to trigger suppression. And together the

experiments I have just described demonstrate the role of suppression in enabling anaphoric reference: Suppression attenuates the interference caused by the activation of other referents.

The Role of Suppression in Cataphoric Reference

Just as anaphoric devices enable reference to previously mentioned concepts, cataphoric devices enable reference to subsequently mentioned concepts. Cataphoric devices include such overt markers as vocally stressing a word in spoken discourse, or bold facing a word in printed text. Presumably speakers and writers mark certain concepts with cataphoric devices because those concepts will play a key role in the text or discourse. Thus, it would behoove listeners and readers if those key concepts had a privileged status in their mental structures.

Masters student Suzanne Shroyer and I (Gernsbacher & Shroyer, 1989) demonstrated that in spoken English, the unstressed, indefinite article *this,* as in *So this man walks into a bar,* as opposed to *So a man walks into this bar,* operates as a cataphoric device. The indefinite *this* is a relative newcomer to English; Wald (1983) dates its use back only to the late 1930s. It occurs almost exclusively in informal spoken dialects rather than formal or written ones, although I have observed personally many of my email pen pals use the indefinite *this* in written email.

Because it is an indefinite article, the indefinite *this* is used to introduce new concepts into discourse. Indeed of the 243 occurrences of the indefinite *this* that Prince (1981) observed in Stud Terkel's book *Working,* 242 introduced a distinctly new concept. More interestingly — particularly with regard to my conjecture that the indefinite *this* operates as a cataphoric device to enable subsequent reference — in 209 of the 243 occurrences of the indefinite *this*, the concept introduced with the indefinite *this* was referred to again. Similarly, when Wright and Givón (1987) recorded 8- and 10-year old children telling one another stories and jokes, they found that when the children introduced concepts with the indefinite *this,* they referred to those concepts an average of 5.32 times in their next 10 clauses. When the children introduced concepts with the indefinite *a/an,* they referred to those concepts only .68 times in their next 10 clauses. These descriptive data suggest that speakers use the indefinite *this* to introduce key concepts. We (Gernsbacher & Shroyer, 1989) tested this proposal experimentally.

We presented spoken narratives to college students, telling them that at some point in each narrative the original narrator would stop talking; when that happened, it was their job to continue. For example, subjects heard the following narrative: *I swear, my friend Vicki, every time we go to a garage sale, she just 'uh, she just goes crazy. I mean like last Saturday we went to one near campus, 'n she just had to buy <u>this/an</u> ashtray, n'man ...* As this example illustrates, the last clause of the part of each narrative that subjects heard introduced a new concept, for example, *ashtray.* We manipulated whether this concept was introduced with the indefinite *this (this ashtray)* or the indefinite *a/an (an ashtray).* When we introduced the concepts with the indefinite *this,* subjects mentioned those concepts considerably more frequently, virtually always within the first clauses that they produced, and usually with less explicit anaphors such as pronouns. In contrast, when we introduced the concepts with the indefinite *a/an,* subjects mentioned the concepts less frequently, and typically with more explicit anaphors such as repeated noun phrases. (Through cross-splicing we ensured that the acoustic properties of the matched narratives and their critical concepts were otherwise identical.) These data demonstrate that concepts marked by cataphoric devices, such as the indefinite this, are more salient in listeners' mental representations.

Recently, postdoctoral fellow Joerg Jescheniak and I (Gernsbacher & Jescheniak, 1995) discovered the role that the cognitive mechanism of suppression plays in enabling this privileged

status: Suppression enables cataphoric reference by attenuating the interference caused by the introduction of other concepts. In this way, a cataphorically marked concept gains that privileged status in comprehenders' mental representations, so that it can be referred to more easily. Consider as an analogy a call for volunteers during which the entire line of candidates steps back, save one. The one candidate who did not step back — who was not suppressed — becomes most accessible for selection.

Subjects in our (Gernsbacher & Jescheniak, 1995) experiments also heard narratives, like the "Vicki going to a garage sale" narrative. We manipulated the indefinite *this* in some experiments, and in other experiments we manipulated a seemingly more powerful cataphoric device, contrastive intonational stress. Using a verification task we measured activation of the experimental concepts. In addition to cataphorically marked concepts being more activated — in other words, enhanced — we also found that cataphorically marked concepts are very resilient to being suppressed by a subsequently introduced concept. For example, in the Vicki going to the garage sale narrative when we introduced a new concept, *vase,* as in *had to buy this/an ashtray, 'n man, then she saw a vase ...* We observed that the previously mentioned concept, *ashtray,* was greatly attenuated in its activation when it was not marked by a cataphoric device. However, when the previously mentioned concept was marked by a cataphoric device, it was just as activated after we introduced a new concept as it was immediately after it was introduced.

Thus, cataphoric devices — the indefinite *this* and contrastive, intonational stress — attenuates the interference caused by introducing other concepts. By attenuating the interference from other concepts, cataphoric devices lead to a privileged status. Furthermore, the two cataphoric devices differ in how powerfully they lead to this privileged status. Figure 4 displays the estimated amount of suppression (the filled bars) and enhancement (the unfilled bars) observed with contrastive stress versus the indefinite *this*. These estimates are based on subjects' reaction times.

Figure 4	Figure 5

As Figure 4 illustrates, suppression and enhancement are more powerfully triggered by contrastive stress than they are by the indefinite *this*. Figure 5 provides a very similar estimate of suppression and enhancement, this time based on subjects' error rates. Again, the figure illustrates that suppression and enhancement are more powerfully triggered by contrastive stress than they are by the indefinite *this*. This difference makes sense: Contrastive stress is considerably more marked; it is a very iconic way of emphasizing a word in spoken discourse — like boldfacing a word in written text. The indefinite *this* is considerably more subtle; many of us are unaware of our informal use of it. Indeed, our undergraduate research assistants, whom we

typically keep blind to our experimental manipulations, were stymied in their attempts to figure out what we were doing in the indefinite *this* experiments. So, like anaphoric devices, the strength of the suppression and enhancement signals triggered by cataphoric devices is a function of the cataphoric devices' markedness.

All the experiments that I have described so far demonstrate the role that suppression plays in attenuating lexical- or concept-level interference. I have also examined the role of suppression in attenuating sentence-level interference.

The Role of Suppression in Syntactic Parsing

Motivated by the adage, *Time flies like an arrow; fruit flies like a banana,* often attributed to Groucho Marks, laboratory coordinator Rachel Robertson and I hypothesized a role that the mechanism of suppression might play in syntactic parsing (Gernsbacher & Robertson, 1997). We proposed that suppression attenuates the interference caused by parsing a previous syntactic form. As the *time flies/fruit flies* example demonstrates, once we have parsed the phrase *time flies* as a noun and verb, it is difficult not to parse the phrase *fruit flies* in the same way. We (Gernsbacher & Robertson, 1997) examined a more stringent type of interference by using phrases such *visiting in-laws,* which can be interpreted either as a plural noun phrase (i.e., people who are related to one's spouse and come to visit) or as a gerundive nominal (i.e., the act of visiting people who are related to one's spouse).

In our experiments (Gernsbacher & Robertson, 1997), we preceded sentences containing phrases like as *visiting in-laws* with sentences that required a similar or conflicting syntactic parse. For example, subjects first read *Washing dishes is a drag,* and then read *Visiting in-laws are, too.* Or subjects first read *Whining students are a drag,* and then read *Visiting in-laws is, too.* The subjects' task was to read each sentence and simply decide whether it was grammatical. We found that subjects were considerably slower and frighteningly less accurate to say that a sentence such as *Visiting in-laws are, too* was grammatical after they read the sentence, *Washing dishes is a drag.* Similarly, subjects were considerably slower and frighteningly less accurate to say that the sentence, *Visiting in-laws is, too* was grammatical after they read the sentence, *Whining students are a drag.* We interpreted these data as suggesting that correctly responding to the second sentence requires attenuating, or suppressing, the interference caused by the syntactic form in the first sentence.

We observed the same effect when we made the second sentences less syntactically dependent on the first sentence, by omitting the ellipses. For example, subjects were again slower and frighteningly less accurate to say that the sentence, *Visiting in-laws are a drag, too* was grammatical after they read the *Washing dishes* sentence. And, subjects were considerably slower and less accurate to say that the sentence *Visiting in-laws is a drag, too* was grammatical after they read the *Whining students* sentence. Furthermore, we observed the same effect when we made the second sentences syntactically independent of the first sentence, and the verb in the first sentence was not even marked for number. For example, subjects were still slower and still frighteningly less accurate to say that the sentence, *Visiting in-laws are a drag* was grammatical after they read the sentence, *Washing dishes can be a bother,* and vice-versa for after they read the sentence *Whining students can be a bother.* This phenomenon underscores the need for suppression to attenuate the interference caused by a previous syntactic form.

The Role of Suppression in Metaphor Interpretation

Rachel Robertson and I, in collaboration with Boaz Keysar (of the University of Chicago) have also explored the role of suppression in metaphor interpretation. According to Glucksberg and Keysar (1990), when we interpret a metaphor such as *Lawyers are sharks,* we should enhance attributes of the metaphor's vehicle (e.g., *sharks)* that are common to the metaphor's topic (e.g., *lawyers).* So, after comprehending the metaphor, *Lawyers are sharks*, we should enhance the facts that sharks are tenacious, fierce, and aggressive, among other attributes. We augmented Glucksberg and Keysar's (1990) theory by proposing that when we interpret a metaphor we also suppress the attributes that are not appropriate to (or concordant with) a metaphorical interpretation. So, for example, when we interpret the metaphor, *Lawyers are sharks,* we might suppress attributes such as sharks being good swimmers, having fins, and living in the ocean. We tested both of these hypotheses by asking subjects to read a statement that might be metaphorical such as *Lawyers are sharks* and then confirm the verity of a property statement such as *Sharks are tenacious.* In our first experiment, we used as a control condition statements that contained the same vehicle but a nonsensical topic, such as *Notebooks are sharks.*

We (Gernsbacher, Keysar, & Robertson, 1995) found striking evidence that interpreting a metaphor such as *Lawyers are sharks* leads to both the enhancement of the attributes that are appropriate to the metaphorical interpretation and the suppression of attributes that are inappropriate to the metaphorical interpretation. For instance, subjects were faster to verify the statement, *Sharks are tenacious* after they read the metaphor, *Lawyers are sharks* than after they read the control statement, *Notebooks are sharks.* This finding supporting the hypothesis that interpreting a metaphor involves enhancing attributes that are appropriate to the metaphorical interpretation. In contrast, subjects were considerably slower to verify the statement, *Sharks are good swimmers* after they read the metaphor, *Lawyers are sharks,* than after they read the control statement, *Notebooks are sharks.* This finding supports the hypothesis that interpreting a metaphor involves suppressing attributes that are inappropriate to the metaphorical interpretation.

In a second experiment, we observed identical results when instead of using a nonsensical statement as a baseline (control), we used a literal statement as a baseline. For example, we presented the literal statement, *Hammerheads are sharks* as a baseline comparison for the metaphorical statement, *Lawyers are sharks.* Again, we found striking evidence to support the hypothesis that interpreting a metaphor leads to both the enhancement of attributes that are appropriate to the metaphorical interpretation and suppression of attributes that are inappropriate to the metaphorical interpretation. For example, again, subjects were faster to verify the statement *Sharks are tenacious* after they read the metaphor *Lawyers are sharks* than after they read the literal statement *Hammerheads are sharks.* And conversely, subjects were again considerably slower to verify the statement *Sharks are good swimmers,* after they read the metaphor *Lawyers are sharks* than after they read the literal statement *Hammerheads are sharks.* Therefore, both experiments demonstrated that interpreting a metaphor involves both enhancing the attributes that are relevant to the metaphorical interpretation and more intriguingly, suppressing the attributes that are not relevant to the metaphorical interpretation.

The Role of Suppression in Revising Inferences

When most of us hear or read that *George became too bored to finish the history book,* we infer that George is reading a very boring book. However if we later hear or read that *George had already spent five years writing it,* we must revise our initially drawn inference because it was inappropriate. Brownell, Potter, Bihrle, and Gardner (1986) found that right-hemisphere damaged patients had a whale of a time revising such inferences. They concluded that right-hemisphere

damaged patients' difficulty arose because they were unable to "let go of" the initial inferences that they drew. Perhaps revising such an inference is difficult because the revision requires suppressing the initially drawn inference. Thus, another role that suppression might play is to attenuate the interference caused by a previously drawn, but erroneous, inference.

We empirically tested this hypothesis, by investigating whether revising such inferences was difficult, not just for right-hemisphere damaged patients but for "normal" college-aged adults. We constructed 40 two-sentence vignettes, similar to the *George became too bored to finish the history book. He had already spent five years writing it* example (other stimuli included *Jeff got a ticket after parking his car. As he headed into the movie theater, he handed the ticket to the usher; Sarah drove frantically all the way to the Emergency Room. She was already running 15 minutes late for her shift that evening; Jack painted the boat a bright red. Then he painted the ocean a deep blue and the sun a bright orange).* We measured how long subjects needed to read the inference-revising second sentence (e.g., *He had already spent five years writing it),* after they read the inference-inviting premise sentence (e.g., *George became too bored to finish the history book).* We compared how long subjects needed to read the inference-revising second sentence (e.g., *He had already spent five years writing it),* after they read the inference-inviting premise sentence or after they read a control premise sentence, which was a sentence that did not invite the inference; indeed it explicitly stated a different situation (e.g., *George became too bored to finish* writing *the history book; Jeff* bought *a ticket after parking his car; Sarah drove frantically all the way to* her job at *the Emergency Room;* In the painting, *Jack painted the boat a bright red).* We found that subjects required substantially longer to read the second sentence (e.g., *He had already spent five years writing it),* after they read the experimental (inference-inviting) premise sentence (e.g., *George became too bored to finish the history book)* than after they read the control (inference-noninviting) premise sentence (e.g., *George became too bored to finish* writing *the history book).* We interpreted subjects' greater latency as reflecting their difficulty in suppressing a previously — but erroneously — drawn inference.

Furthermore, we found that members of a particular subject group — a group which we have previously identified to have difficulty quickly employing suppression (as I shall describe below) — were substantially slower to reject a test word that was related to the erroneously drawn inference, even after they read the inference-revising second sentence. For example, members of this group of subjects took longer to reject the test word *READ* even after they read the inference-revising sentence *He had already spent five years writing it.* Members of this subject group were less-skilled comprehenders, and this finding leads me to the last section of this introduction, the role of suppression in general comprehension skill.

The Role of Suppression in Comprehension Skill

A few years ago, Kathy Varner, Mark Faust and I presented evidence in support of a construct we called "General Comprehension Skill" (Gernsbacher, Varner, & Faust, 1990). Briefly put, we found that adults' skill in comprehending written language was highly correlated with their skill in comprehending spoken language, and both skills were highly correlated with comprehending nonverbal picture stories. We also found a critical characteristic of less-skilled adult comprehenders: Less-skilled adult comprehenders are less able to suppress quickly the inappropriate meanings of homonyms.

We (Gernsbacher et al., 1990) discovered this critical characteristic in the following way: We selected 64 more- versus less-skilled University of Oregon undergraduates on the basis of their performance on our Multi-Media Comprehension Battery (Gernsbacher & Varner, 1988). This

battery tests reading, listening, and picture story comprehension. We drew the more-skilled comprehenders from the upper third of a distribution of 270 subjects and the less-skilled comprehenders from the bottom third. We invited these more- and less-skilled subjects to return to our lab (which was no easy feat, as the less-skilled subjects did not have that much fun the first time they were there). When the subjects returned they read short sentences, and following each sentence, they were shown a test word. Their task was to decide quickly whether the test word fit the meaning of the sentence that they just read. On experimental trials, the final-word of the sentence was a homonym, such as *spade* as in *He dug with the spade.* The test word on these trials was related to a meaning of that homonym, but not the meaning implied by the sentence, for example, *ACE*. We compared how rapidly the more- vs. less-skilled comprehenders could reject a test word that was related to the inappropriate meaning with how rapidly they could reject the same test word after reading a control sentence, for example, *He dug with the shovel.* The more time subjects took to reject *ACE* following the spade- versus shovel-sentence, the more interference they were experiencing from the inappropriate meaning. We measured this interference immediately (100 ms) after subjects finished reading the sentences and after an 850 ms delay. Figure 6 displays our data, presented as interference scores; the more skilled-comprehenders are represented by the hashed lines and the less-skilled comprehenders are represented by the unfilled bars.

| | Figure 6 | | Figure 7 |

As Figure 6 illustrates, immediately after both the more- and less-skilled comprehenders read the homonyms, both groups experienced a reliable amount of interference, and indeed, the two

groups did not differ in the amount of interference they experienced at the immediate test point. In contrast, after the delay, the more-skilled comprehenders were no longer experiencing a reliable interference, suggesting that they had successfully suppressed the inappropriate meanings of the homonyms. But for the less-skilled comprehenders, they experienced the same amount of interference after they delay as they experienced immediately, suggesting that they were less able to quickly suppress the inappropriate meanings of the homonyms.

This pattern has been replicated by our colleagues around the world with whom we have shared our stimuli. For example, Leslie Twilley and Peter Dixon at the University of Alberta replicated this pattern using our measure of comprehension skill. Harvey Shulman at Ohio State replicated this pattern testing subjects who scored in the top versus bottom half of the ACT verbal test, but not math test. Francesca Pazzaglia and her colleagues at the University of Padova replicated this pattern using Italian homonyms (and Italian subjects). Natasha Todorov at Macquarie University replicated this pattern with seventh-graders selected according to their Nelson-Denny reading scores. In his dissertation Bob Crane at Washington State replicated this pattern testing university students, with small versus large reading spans. And we replicated this pattern testing United States Air Force recruits. Thus, this pattern replicates with Canadians, Italians, Australians, Buckeyes, and the U.S. military.

Furthermore, this pattern occurs when comparing members of other populations who hypothetically suffer from less efficient suppression with members of populations who are hypothesized to have more efficient suppression. For example, using our task and stimulus materials, McDowd and Sundry at the University of Southern California found that healthy elderly subjects showed less efficient suppression compared with college-aged subjects. Elizabeth Schaunessy found that children diagnosed with attention deficit disorder showed less efficient suppression compared with children not diagnosed with attention deficit disorder. Mark Faust, David Balota, Janet Duckek and I found that patients with severe senile dementia of the Alzheimer's type showed extraordinarily inefficient suppression compared with patients with only moderate dementia compared with healthy age-matched controls (Faust et al., 1997). Indeed, our dementia data show something of a dosing effect: The more severe the dementia, the more inefficient the suppression.

In all the data that I have reviewed, all the subjects showed initial interference from the inappropriate meanings, which I believe is crucial to demonstrate, but the members of the population hypothesized to suffer from less-efficient suppression showed continued interference from the inappropriate meanings. Mark Faust and I (Gernsbacher & Faust, 1991b) also observed the same pattern when we examined how quickly less- versus more skilled comprehenders could reject test words related to the incorrect forms of homophones, for instance, how quickly they could reject the test word *CALM* following the sentence, *He had lots of patients,* versus *He had lots of students.* Figure 7 displays our data (from US Air Force recruits). Prior to collecting these data, we conducted a pilot experiment to ensure that members of this population did know the correct spelling of these homophones.

We (Gernsbacher & Faust, 1991b) also discovered that less- versus more-skilled comprehenders are not less able to reject the contextually inappropriate meanings of homonyms just because they do not know what is appropriate. We observed that less-skilled comprehenders perform equally as well as more-skilled comprehenders when the task is to accept the appropriate meaning of a homonym, for example, when their task is to correctly say "yes" that the test word *ACE* is related to the sentence, *He dealt the spade.*

Very recently, Rachel Robertson and I (Gernsbacher & Robertson, 1995) replicated our tried and true finding that less-skilled comprehenders are worse than more-skilled comprehenders

when the task is to reject a test word that is related to the inappropriate meaning. For example, less-skilled comprehenders are slower to reject the test word *ACE* after reading the sentence, *He dug with the spade.* Presumably this is because less-skilled comprehenders are less able to suppress the activation of the inappropriate meanings. We (Gernsbacher & Robertson, 1995) also replicated the finding I just mentioned, namely that less- and more-skilled comprehenders do not differ when the task is to accept the appropriate meaning. For example, less-skilled comprehenders are just as fast as more-skilled comprehenders in accepting the test word *ACE* after reading the sentence, *He dealt the spade.* Again, this suggests that less-skilled comprehenders' difficulty in rejecting inappropriate meanings is *not* because they do not know what is appropriate.

However, we (Gernsbacher & Robertson, 1995) also created a task in which the goal was to say "yes" to a meaning that was inappropriate, somewhat like what one needs to do to understand a pun. And we again found that less-skilled comprehenders were worse than more-skilled comprehenders. For example, less-skilled comprehenders were slower to accept the test word *ACE* after reading the sentence *He dug with the spade,* perhaps because this task — accepting an inappropriate meaning — requires suppressing the appropriate meaning (recall how difficult it was to understand the pun, *Two men walk into a bar and the third man ducks.* It is as though to understand the "metal bar" meaning, one needs to suppress the "tavern" meaning).

To summarize, I have suggested that the cognitive mechanism of suppression plays a crucial role in many language comprehension phenomena. During lexical access, the mechanism of suppression attenuates the interference caused by the activation of other lexical information, such as the inappropriate meanings of homonyms. During anaphoric reference, the mechanism of suppression attenuates the interference caused by the activation of other potential referents. In this way, the referent to which the anaphor does refer becomes the most activated concept. Moreover, the strength of the suppression is a function of the markedness of the anaphoric device: More marked anaphors such as repeated proper names immediately lead to suppression; less marked anaphors such as pronouns take longer to enact suppression. During cataphoric reference, the mechanism of suppression attenuates the interference caused by the introduction of other concepts. In this way, a cataphorically marked concept gains a privileged status in comprehenders' mental representations. The more marked the cataphoric device is, the less interference is caused by the introduction of another concept. More marked cataphoric devices such as spoken stress protect their concepts more than less marked cataphoric devices such as the indefinite *this.* During syntactic parsing, the mechanism of suppression attenuates the interference caused by a previous syntactic form. During metaphor comprehension, the mechanism of suppression attenuates the interference caused by a literal interpretation. During inferencing, the mechanism of suppression attenuates the interference caused by an initial but inappropriate inference. Thus, my previous research has demonstrated the crucial role that suppression — and by that I mean a general, cognitive mechanism that attenuates interference — plays in many facets of language comprehension.

Investigations of the Strategic Control of Suppression

The research I recently conducted for the Army Research Institute extended the research I conducted during the previous funding period. During the previous funding period, I answered the following questions: *1) Is the mechanism of suppression under strategic control? 2) If the mechanism of suppression is strategically controlled, can persons be taught to suppress competing information?* Both questions were answered by a series of laboratory experiments.

The experiments that answered the first question explored the relative automaticity versus strategic control involved in suppressing irrelevant, inappropriate, and to-be-ignored lexical, pictorial, and grammatical information.

The literature on attention differentiates automatic mental activity and cognitive processes from those that are more controlled, perhaps strategic (Keele & Neill, 1978; Posner & Snyder, 1975a, 1975b). Automatic versus strategic mechanisms can be differentiated in the laboratory by manipulating the probability of a particular type of trial occurring within an experiment. The logic of a probability manipulation is this: If a certain type of experimental trial occurs only rarely, subjects are unlikely to adopt a strategy for that type of trial. But if a certain type of trial occurs frequently, subjects are likely to adopt a strategy for successfully doing that type of trial.

Consider the following experimental task: Subjects see pairs of letter strings, appearing side by side (e.g., *DORTZ BLAUGH*). The subjects' task is to decide whether each member of the pair is a word. On some trials, both members are words, and on some of the trials in which both members are words, the two words are semantically related, for example, *BREAD BUTTER*. A classic finding is that the second letter string is recognized more rapidly when it appears in a pair of related words; for example, *BUTTER* is recognized more rapidly when it is appears in the related-word pair *BREAD BUTTER* than when it appears in the unrelated-word pair *NURSE BUTTER* (Meyer & Schvanaveldt, 1971). Now consider the following manipulation: In one condition, only 1/8 of the word pairs are related *(BREAD BUTTER)*, and the majority (7/8) are unrelated *(NURSE BUTTER)*; in another condition 1/2 are related, and 1/2 are unrelated; and in a third condition, the majority (7/8) of the word pairs are related, and only 1/8 are unrelated. With this manipulation, subjects recognize the second word of the pair more rapidly if the pair is related (just as other experiments have shown), and the advantage of the relatedness between the two words in a pair is a function of the probability of a related word pair appearing in the experiment. When only 1/8 of the word pairs in the experiment are related, the advantage is smallest; when 7/8 of the word pairs are related, the advantage is largest. Presumably, the high probability of related pairs encourages subjects to adopt a strategy for capitalizing on the words' relations (Tweedy, Lapinski, & Schvanaveldt, 1977).

In other experiments, subjects also adopt beneficial strategies in response to a high probability of a certain type of experimental trial. For instance, in experiments in which subjects perform a letter matching task, subjects are shown pairs of letters, and they decide rapidly whether the members of the pair match (either physically, e.g., *A* and *A,* or in name, e.g., *a* and *A).* In Posner and Snyder's (1975a) experiment, the letter pairs were preceded by three types of cues: an informative cue, which was one of the letters of the pair (e.g., the cue was *A,* and the pair was *AA),* a neutral cue (a plus sign), or an uninformative cue, which was a letter that did not match either member of the pair (e.g., the cue was *B,* and the pair was *AA).* Posner and Snyder (1975a) varied the probability of the cue being informative. It was informative on 20%, 50%, or 80% of the trials. Subjects were fastest when the cue was informative, and when the informative cue occurred 80% of the time. Presumably with a high probability of informative cues, subjects adopted a strategy for taking advantage of the informative cues.

However, subjects do not always adopt a strategy, even when there is a high probability of a particular type of trial. Subjects adopt a strategy *only* if the cognitive mechanism tapped by that type of trial is under the subjects' strategic control. If the cognitive mechanism tapped by the highly probable type of experimental trial is not under the subjects' strategic control, then the probability manipulation is ineffective. For instance, in an experiment in which subjects have to decide whether each member of a pair of letter strings are words, subjects typically adopt a beneficial strategy for dealing with word-word pairs, when there is a high probability of that the

two words are semantically related (e.g., *BREAD BUTTER*). However, subjects adopt a strategy only if they have enough time to process the first word of the pair; without adequate time for processing the first word, a 1/8 vs. 1/2 vs. 7/8 ratio of related- to unrelated-word pairs has no effect (den Heyer, Briand & Dannenbring, 1983).

Consider another situation in which subjects cannot employ an adaptive strategy. In an experiment conducted by Simpson and Burgess (1985), subjects first read an ambiguous prime word, such as *BANK*. After 750 ms, each prime word disappeared, and the subjects saw a test word, such as *MONEY*. The subjects made a lexical decision to each test word. On some trials, the test words were related to the most-frequent meaning of the ambiguous prime words. For example, *MONEY* is related to the most-frequent meaning of *BANK*. On other trials, the test words were related to a less-frequent meaning of the ambiguous prime words. For example, *RIVER* is related to a less-frequent meaning of the ambiguous prime word *BANK*.

Simpson and Burgess (1985) measured how rapidly subjects responded to the test words *(MONEY* or *RIVER)* when the prime words were ambiguous *(BANK)* versus unambiguous (e.g., *RIDDLE)*. Simpson and Burgess (1985) also manipulated the probability that the test words were related to the less-versus more-frequent meanings of the ambiguous prime words. In one condition, the test words were related to the less-frequent meanings on the majority, 80%, of the trials, and on only 20% of the trials were the test words related to the more-frequent meanings. In another condition, the test words were related to the less- versus more-frequent meanings on an equal number of the trials (50%); and in a third condition, the test words were related to the less-frequent meanings on only 20% of the trials, and they were related to the more-frequent meanings on 80% of the trials.

Simpson and Burgess (1985) found that the probability manipulation was ineffective. Regardless of the probability that the test words would be related to the less- versus more-frequent meanings, subjects recognized (made lexical decisions to) the more-frequent meanings *(MONEY)* more rapidly than they recognized the less-frequent meanings *(RIVER)*. Thus, even when the test words were related to the less-frequent meanings on 80% of the trials, subjects still recognized the more-frequent meanings more rapidly than they recognized the less-frequent meanings (just as they did when the test words had an equal probability of being related to the less- versus more-frequent meanings). In fact, in a fourth condition, subjects were informed that many of the prime words would be ambiguous and that 80% of the test words would be related to those prime words' less-frequent meanings. But even with this informative warning, subjects still did not recognize the less-frequent meanings more rapidly than they recognized the more-frequent meanings. These data suggest that subjects cannot adopt a strategy to improve their recognition of the less-frequent meanings of homonyms. Perhaps the ability to rapidly recognize the less-frequent meanings of homonyms is driven by a cognitive mechanism that is not under subjects' strategic control.

Suppression During Lexical Access Is Under Strategic Control

In one set of experiments that I conducted during the recent funding period, I used a probability manipulation to investigate whether the cognitive mechanism of suppression is under subjects' strategic control. In one experiment I investigated whether the suppression that is employed during the process of lexical access (i.e., the attenuation of interference caused by contextually-inappropriate meanings of homonyms) was under strategic control. This experiment was based on an experiment reported in Gernsbacher et al. (1990), in which I demonstrated that correctly understanding a sentence that contains a homonym requires actively suppressing the

meanings of that homonym that are not implied by the sentence's context. For example, correctly understanding the sentence, *He dug with the spade,* requires suppressing the meaning of *spade* that is associated with playing cards.

In my previous experiment (Gernsbacher et al., 1990; Experiment 4), subjects read short sentences; after reading each sentence, the subjects saw a test word. Their task was to verify whether the test word fit the meaning of the sentence they just read. On 80 trials, the test word did indeed fit the meaning of the sentence, but I was more interested in the 80 trials in which the test word did <u>not</u> fit the sentence. On 40 of those trials, the last word of the sentence was a homonym, for example, *He dug with the spade.* The test word on these trials was a meaning of the homonym that was inappropriate to the meaning of the sentence, for example, *ACE.* (Unlike the Simpson and Burgess, 1985, experiments all the homonyms were words for which both the appropriate and inappropriate meanings were equal in their frequency of interpretation, when those homonyms were presented in isolation. That is, there were no high- versus low-frequency meanings.)

I measured how long subjects took to reject a test word like *ACE* after reading a sentence like, *He dug with the spade.* And I compared that latency with how long subjects took to reject *ACE* after reading the same sentence but with the last word replaced by a unambiguous word, for example, *He dug with the shovel.* This comparison indicated how activated the inappropriate meaning of the homonym was; the more time subjects took to reject *ACE* after the *spade-* versus the *shovel-*sentence, the more activated the inappropriate meaning must have been. I found that immediately (100 ms) after subjects read the homonyms, the inappropriate meanings were highly activated. However, after a one-second delay, the inappropriate meanings were no longer activated (for subjects who were above-average in their general comprehension skill). I concluded that the inappropriate meanings were no longer highly activated after the one-second delay, because the subjects had actively suppressed those inappropriate meanings.

I discovered that the suppression of inappropriate meanings is under strategic control in the following way. Subjects read short sentences, and after each sentence, they saw a test word. Their task was to verify whether the test word fit the meaning of the sentence they just read. On 80 trials, the test word did indeed fit the sentence, but I was more interested in the 80 trials in which the test word does <u>not</u> fit the sentence. In these 80 trials, the sentence-final word was either a homonym (e.g., *spade)* or an unambiguous word (e.g., *shovel).* I manipulated the probability that the sentence-final word was a homonym versus an unambiguous control word. In the high-probability condition, the sentence-final word was a homonym on the majority, 75%, of the trials and an unambiguous word on only 25% of the trials. In the low-probability condition, the sentence-final word was a homonym on only 25% of the trials and an unambiguous word on the majority, 75%, of the trials.

The test word on both types of trials was related to a meaning of the homonym that was inappropriate to the context, for example, *ACE.* Rejecting a test word like *ACE* following a homonym like *spade* requires suppressing the inappropriate meaning. Rejecting *ACE* following an unambiguous word like *shovel* does not require this suppression. If the suppression of inappropriate meanings during lexical access is under subjects' strategic control, then subjects should be more likely to suppress the contextually inappropriate meanings in the high-probability condition than in the low-probability condition. And that is exactly what I found.

Suppression During Scene Recognition Is Under Strategic Control

In Gernsbacher and Faust (1991a), I demonstrated that accurately viewing a visual scene requires actively suppressing objects that typically occur in such scenes, but are not present in the actual scene being viewed. I am not alone in claiming that accurate perception of visual scenes often involves the cognitive mechanism of suppression. For instance, Biederman writes about the difficulty in "suppressing the interpretations of visual arrays that comprise scenes" (Biederman, Bickle, Teitelbaum, & Klatsky, 1988, p. 456). This difficulty is manifested in the following experimental phenomenon: Subjects are more likely to incorrectly report that a tractor was present in a scene of a farm than in a scene of a kitchen (Biederman, Glass, & Stacy, 1973; Biederman, Mezzanotte, & Rabinowitz, 1982; Biederman, Teitelbaum, & Mezzanotte, 1983; Palmer, 1975).

To successfully comprehend a scene, I have claimed that observers must suppress these typical-but-absent objects, just as readers and listeners must suppress the inappropriate meanings of homonyms. In another experiment that I conducted during the recent funding period I investigated whether the suppression of typical-but-absent objects is under strategic control. This experiment was modeled after Gernsbacher and Faust (1991a), Experiment 2. In Gernsbacher and Faust (1991a), Experiment 2, subjects viewed arrays of objects that were typical of a particular scene, such as six objects from a farm scene. After viewing each array, subjects saw the name of another object. Their task was to verify whether the object named by the test word had been present in the array they just viewed. On 80 trials, the object named by the test word had been present in the array the subjects viewed, but in 80 trials, the object named by the test word was not present in the array the subjects viewed. I was interested in the trials in which the object was absent.

On half of those trials, the objects in the array were typical of a particular scene, for instance, objects that typically occur in a farm scene (e.g., *farmer, pitchfork, pig, barn, chicken, ear of corn)*. On these trials, the test word was the name of an object that also typically occurs in this type of scene (e.g., *TRACTOR),* but that object had not been present in the scenic array that the subjects just viewed. I measured how long subjects took to reject a test word like *TRACTOR* after viewing an array comprising other objects that belong to a farm scene, and I compared that latency with how long they took to reject *TRACTOR* after viewing an array of objects that typically belong to another scene, for instance, objects belonging to a kitchen scene. This comparison indicated how activated the typical-but-absent object was; the longer subjects took to reject *TRACTOR* after viewing the typical *(farm)* array versus the atypical *(kitchen)* array, the more activated the typical-but-absent object must have been. I found that immediately (50 ms) after subjects viewed the arrays, the typical-but-absent objects were highly activated (a finding that replicates Biederman's work). However, after a one-second delay, the typical-but-absent objects were no longer activated (for subjects who were above-average in their general comprehension skill). I concluded that the typical-but-absent objects were no longer activated after the one-second delay because the subjects had actively suppressed them.

I discovered that the suppression of typical-but-absent objects is under strategic control in the following way. Subjects viewed arrays of objects that were typical of a particular scene, such as objects from a farm scene. After viewing each array, subjects indicated whether an object named by a test word was present in the array they just viewed. On 80 trials, the object named by the test word had been present, but on 80 it had not. I was interested in those 80 trials in which the object named by the test word was absent. I manipulated the probability that the array was typical versus atypical of the (absent) test object. In the high-probability condition, the array was typical on the majority, 75%, of the trials and atypical on only 25% of the trials. In the low-probability condition, the array was typical on only 25% of the trials and atypical on the

majority, 75%, of the trials. The test word on both types of trials was the same: the name of the absent object, for example, *TRACTOR*. Rejecting a test word like *TRACTOR* following an array of farm objects requires suppressing the typical-but-absent object. Rejecting *TRACTOR* following an array of kitchen objects does not require this suppression. If the suppression of typical-but-absent objects is under subjects' strategic control, then subjects should be more likely to suppress the typical-but-absent objects in the high-probability condition than in the low-probability condition. And that is exactly what I found.

Suppression During Syntactic Parsing Is Under Strategic Control

In Gernsbacher and Robertson (1997) we discovered a compelling phenomenon that suggests that the cognitive mechanism of suppression plays a crucial role in syntactic parsing (which is the process of decoding the grammatical form of a sentence so that the reader or listener can understand the sentence's meaning). Subjects in these experiments read pairs of sentences. The subjects were told that the two sentences within each pair were unrelated, and they should attempt to understand each sentence independently from the other. The subjects' overt task was to decide whether each sentence was grammatical. For example, the sentence, *Washing clothes is often a bother,* is grammatical, whereas the sentence, *Washing clothes are often a bother,* is not grammatical. As another example, the sentence, *Whining students are often a bother,* is grammatical, whereas the sentence, *Whining students is often a bother,* is not grammatical. In one-fourth of the sentence pairs, neither the first nor second sentence was grammatical. These pairs served as filler trials. In another fourth of the sentence pairs, the first sentence was not grammatical, but the second sentence was. Again, these pairs served as filler trials. In another fourth of the sentence pairs, the first sentence was grammatical, but the second sentence was not. And again, these served as filler trials. The remaining 64 sentence pairs comprised the experiment trials, and on experimental trials, both sentences of the pair were grammatical.

All sentences were of the syntactic form *Noun phrase is/are adjective* (e.g., *Washing clothes is a bother,* or *Whining students are a bother)*. The critical characteristic of the experimental trials was the second sentence of each experimental sentence pair. All the second sentences contained a head noun phrase that if presented in isolation would be ambiguous. For example, the phrase *Visiting in-laws* can be interpreted either as a plural noun, as in *Visiting in-laws are a bother,* or as a gerundive nominal, as in *Visiting in-laws is a bother.* On half the experimental trials, the head noun phrase in the second sentence of each pair should have been interpreted as a plural noun, because it was followed by the plural verb *are (Visiting in-laws are a bother).* On the other half of the experimental trials, the head noun phrase in the second sentence should have been interpreted as a gerundive nominal, because it was followed by the singular verb *is (Visiting in-laws is a bother).* Thus, the verb in the sentence *(is* or *are)* clearly indicated whether the phrase should have been interpreted as a plural noun *(Visiting in-laws are a bother)* or as a gerundive nominal *(Visiting in-laws is a bother).* Unknown to the subjects, the experimental sentence pairs were constructed in the following way. The first sentence of each experimental sentence pair contained an unambiguous noun phrase that could be interpreted as only a plural noun or as only a gerundive nominal (e.g., *washing dishes* or *whining students).* Half the first sentence noun phrases matched the syntactic form of the second sentence noun phrases, and half mismatched. An example of a matching sentence pair is *Washing dishes is a bother. Visiting in-laws is a bother, too.* An example of a mismatching sentence pair is *Whining students are a bother. Visiting in-laws is a bother, too.*

We found that subjects were considerably slower and less accurate to decide that the second sentences were grammatical when those second sentences mismatched (rather than matched) the first sentences' syntactic form. For example, subjects were considerably slower and less accurate to decide that the second sentence in the following pair was grammatical: *Washing dishes is a bother. Visiting in-laws are a bother, too.* Similarly, subjects were considerably slower and less accurate to decide that the second sentence in the following pair was grammatical: *Whining students are a bother. Visiting in-laws is a bother, too.* Indeed, when the second sentences mismatched the first sentences in syntactic form, subjects were barely above chance at correctly responding that the second sentences was grammatical. These data suggest that correctly interpreting a sentence's syntactic form (i.e., parsing a sentence) requires suppressing a previously presented syntactic form. (Note, however, that this result is not due to the presence of the mismatching verb *is* or *are*. In one of the experiments in Gernsbacher & Robertson, 1997, the first sentences did not contain the verb *is* or *are*. Instead, they contained the verb *can be*, as in *Washing dishes can be a bother. Visiting in-laws are a bother, too,* or *Whining students can be a bother. Visiting in-laws is a bother, too.* The results of this experiment were identical to those in which the first sentences contained the verb *is* or *are*. Thus, subjects direct interpretation of the first sentences' noun phrases affected their interpretation of the second sentences' noun phrases.)

I discovered that the suppression of prior syntactic form is under subjects' strategic control in the following way. Subjects read sentence pairs, and their task was to decide whether each sentence of the pair was grammatical. As in our previous experiments, subjects were told that the two sentences comprising a pair were unrelated, and, therefore, the two sentences should be interpreted independently. Also, as in our previous experiments, there was an equal number of filler sentence pairs in which the first, the second, or both sentences of a pair were not grammatical. However, as in our previous experiments, both sentences comprising the experimental sentence pairs were grammatical. The head noun phrase in the first sentence of each experimental sentence pair was unambiguous, both in isolation and in its sentence context (e.g., *Washing dishes,* or *Whining students).* The head noun phrase in the second sentence of each experimental sentence pair was ambiguous, if presented in isolation. However, on some trials the noun phrases in the second sentence was clearly disambiguated by either the verb *is* or *are*. On these trials the second sentence overtly mismatched the syntactic form of the first sentence. On other trials, the noun phrase in the second sentence was not disambiguated because the verb was the modal + infinitive verb *can be,* as in *Visiting in-laws **can be** a bother.* The modal + infinitive verb *can be* does not indicate whether the noun phrase is a plural noun or a gerundive nominal. Presenting the second sentences with the verb *can be* provided a baseline from which I could more accurately assess suppression. In the high-probability condition, the second sentence mismatched the syntactic form of the first sentence on the majority, 80%, of the trials, and was left ambiguous on only 20% of the trials. In the low-probability condition, the second sentence mismatched the syntactic form of the first sentence on only 20% of the trials, and was left ambiguous on the 80% of the trials. If the suppression of prior syntactic form is under subjects' strategic control, then subjects should respond more accurately and more rapidly in the high-probability condition than in the low-probability condition. And again, this is exactly what I found.

Thus, during the recent funding period I discovered that the mechanism of suppression during lexical access, scene recognition, and syntactic parsing is under strategic control. During the recent funding period I also discovered that persons can be taught to suppress competing information. The experiments that lead to this discovery were based on a suggestion by Wegner (1989; 1992). Wegner suggested that persons can strategically suppress competing information

(what he calls "unwanted thoughts") by purposely focusing on — rather than ignoring — the unwanted thoughts. So, instead of trying to disregard the unwanted thoughts, Wegner suggests that people consciously examine them. By this he means bringing the unwanted thoughts into conscious awareness and examining why the thoughts are unwanted. This suggestion contrasts with most people's intuitions about how to get rid of unwanted thoughts; most people intuitively believe that by forgetting about things one doesn't want to think about, they will "go away." I applied Wegner's suggestions to the strategic control of the cognitive mechanism of suppression during lexical access, scene recognition, and syntactic parsing and discovered that subjects can learn to suppress the contextually inappropriate meanings of homonyms during lexical access; subjects can learn to suppress typical-but-absent objects during scene recognition; and subjects can learn to suppress a previous grammatical form during syntactic parsing.

Suppression During Lexical Access Can Be Learned

As described earlier, correctly rejecting the test word *ACE* following the sentence, *He dug with the spade* requires suppressing the contextually-inappropriate meaning of the word, *spade*. During the previous funding period I discovered that subjects can learn to suppress these contextually-inappropriate meanings. Half the subjects underwent a training session prior to performing the experimental task. The experimental task involved reading a short sentence, and after reading each sentence, verifying whether a test word fits the meaning of the sentence just read. Trials that required suppression were those in which the last word of the sentence was a homonym (e.g., *He dug with the spade),* and the test word was a meaning of the homonym that was inappropriate to the meaning of the sentence (e.g., *ACE).*

During the training session, subjects were presented with 40 training sentences. The subjects were told explicitly that some of the 40 sentences would end in an homonym (i.e., a word that has several unrelated meanings). The subjects were given several examples (e.g., *He dug with the spade)*. During the training session, subjects were also presented with a test word following each of the 40 training sentences. The subjects were told that some of the test words following each training sentence were related to a meaning of the sentence-final homonyms, but not the meaning of the homonyms that is appropriate, given the sentence context (e.g., *ACE)*. For each of the 40 sentences encountered during the training phase, subjects were required to overtly classify test words that are related to the inappropriate meanings, for example, classify that *ACE* following the sentence, *He dug with the spade*, is related to an inappropriate meaning of *spade.* The subjects were also required to state why the test word was inappropriate (e.g., stating something such as, "ACE refers to a playing card, but the type of spade suggested in the sentence is a garden tool"). Each of the 40 training sentences were presented four times during the training session. Following the training session, the subjects performed the reaction time task described above (i.e., read a sentence such as *He dug with the spade* and rapidly decided whether a test word such as *ACE* was fit the meaning of the sentence). Half the experimental items were those on which the subjects had been trained, and half were novel. In this way, I was able to compare whether the training transfers to a novel set of stimuli. Furthermore, by comparing the performance of the subjects who underwent the training with the performance of subjects who did not, I was able to assess the efficacy of the training intervention. The training was highly successful.

Suppression During Scene Recognition Can Be Learned

As described earlier, correctly rejecting the test item *TRACTOR* as having not been present in an array of objects typically found in a farm scene requires suppression. In the previous funding period I discovered that subjects can learn to suppress these typical-but-absent members of scenic arrays. Half the subjects tested underwent a training session prior to performing the experimental task. The experimental task involved viewing an array of objects, and after viewing each array, verifying whether a test object was present. Trials that required suppression were those in which the test object was typically found in a scene that comprises the objects that were present in the array (e.g., a *farmer, pitchfork, pig, barn,* and *chicken* typically occur in a farm scene). During the training session, subjects were presented with 40 scenic arrays. The subjects were told that the scenic arrays would comprise objects that are typically found in the same scene. The subjects were also told that some of the test words would be the names of objects that also typically occur in one of those scenes, but the object would not have appeared in the scenic array they viewed. Subjects were required to identify overtly the type of scene for which the items in the array were typical (e.g., say something such as "these are all objects that are found in a farm scene"). The subjects were also required to pronounce the name of the test object aloud and to overtly identify whether the named test object appeared in the scenic array. After overtly identifying whether the test object was present, the subjects were required to identify whether the test object would typically be found in that type of scene (e.g., "Yes, a tractor would typically be found in a farm scene"). Then, the subjects were required to recall the objects that were present in the array, and again confirm whether the test object had been present (e.g., "there was a farmer, pitchfork, pig, barn, chicken, and an ear of corn, but there was <u>not</u> a tractor").

Following the training session, the subjects performed the reaction time task described above (i.e., they were presented with a scenic array and then the name of a test object and their task was to rapidly verify whether the test object had been presented in the scenic array). Half the stimuli were those on which the subjects had been trained, and half were novel. In this way, I was able to compare whether the training transferred to a novel set of stimuli. Furthermore, by comparing the performance of the subjects who underwent the training with the performance of subjects who did not undergo the training, I was able to assess the efficacy of the training intervention. The training intervention was highly successful.

Suppression During Syntactic Parsing Can Be Learned

As described earlier, correctly deciding that the sentence, *Visiting in-laws are a bother* was grammatical was more difficult if subjects had previously read the sentence, *Washing clothes is a bother*, than if subjects had previously read the sentence, *Whining students are a bother*. I concluded that correctly accepting the sentence, *Visiting in-laws are a bother*, requires suppressing the syntactic frame computed when reading the first sentence. During the previous funding period I discovered that subjects can learn to suppress the syntactically incorrect parse of a noun phrase. Half the subjects tested underwent a training session prior to performing the experimental task. The experimental task involved reading pairs of sentences, and deciding after reading each sentence whether it was grammatical. Trials which required suppression were those in which the interpretation of the noun phrase in the second sentence (e.g., *Visiting in-laws is*) differed from the interpretation of the noun phrase in the first sentence (e.g., *Whining students are* versus *Washing clothes is*).

During the training session, subjects were presented with 32 pairs of training sentences. The subjects were told explicitly that the noun phrases in the second sentences would be ambiguous, if they were presented in isolation. The subjects were given examples. The subjects were also

told that some of the first sentences would have a different syntactic form than the second sentences. And again, they were given examples. Subjects were required to identify overtly the syntactic form of the second sentences. For example, to the training sentence, *Visiting in-laws is a bother,* subjects were to say something like: "This sentence implies that the action of visiting in-laws is a bother, and yes, there is an action that we refer to as visiting people who are called in-laws". Or as another example, to the training sentence, *Visiting in-laws are a bother,* subjects might say, "This sentence implies that things — the in-laws who are visiting — are a bother, and yes, those things [the in-laws] can be described as performing that action."

Following the training session, the subjects performed the reaction time task described above. Half the stimuli presented during the reaction time task were stimuli on which the subjects were trained, and half were novel. Thus, I was able to compare whether the training transferred to a novel set of stimuli. And again, by comparing the performance of the subjects who underwent the training with the performance of subjects who did not undergo the training, I was able to assess the efficacy of the training intervention. And again, the training intervention was highly successful.

Transferring the Skill of Suppression

The research that I conducted for the Army Research Institute during the recent funding period extended the two sets of discoveries described above that deal specifically with the strategic control and training of suppression skill. The recent research answered the following question: *Does training in suppression of one type of competing information transfer to skill in suppressing another type of competing information?* This question was answered through a series of laboratory experiments, the results of which informed us about the transferability of training of the crucial skill of suppression.

During the previous funding period, I discovered that training in assessing competing lexical, pictorial, and grammatical information improved subjects' ability to suppress competing lexical, pictorial, and grammatical information. These training interventions were "domain specific," that is, subjects who were taught to assess the nature of competing lexical information (i.e., the inappropriate meanings of homonyms) improved their ability to suppress competing lexical information; subjects who were taught to assess the nature of competing pictorial information (i.e., the typical-but-absent members of scenic arrays) improved their ability to suppress competing pictorial information; and subjects who were taught to assess the nature of competing grammatical information (i.e., the incorrect parse of syntactic phrases) improved their ability to suppress competing grammatical information.

Does this training transfer across domains? That is, does training in assessing competing lexical information improve performance in suppressing competing pictorial or grammatical information? Does training in assessing competing pictorial information improve performance in suppressing competing lexical or grammatical information? And does training in assessing competing grammatical information improve performance in suppressing competing lexical or pictorial information? During the recently completed funding period, I was able to answer the first of these questions.

I investigated whether training in assessing competing lexical information (the inappropriate meanings of homonyms) improves subjects' ability to suppress competing pictorial information (the typical-but-absent members of scenic arrays). Half the subjects underwent a training session prior to performing the experimental task. During the training session, subjects were presented with 40 training sentences. The subjects were told explicitly that some of the 40 sentences would

end in a homonym (i.e., a word that has several unrelated meanings). The subjects were given several examples (e.g., *He dug with the spade)*. During the training session, subjects were also presented with a test word following each of the 40 training sentences. The subjects were told that some of the test words following each training sentence will be related to a meaning of the sentence-final homonyms, but not the meaning of the homonyms that is appropriate, given the sentence context (e.g., *ACE)*. For each of the 40 sentences encountered during the training phase, subjects were required to overtly classify test words that were related to the inappropriate meanings, for example, classify that *ACE* following the sentence, *He dug with the spade*, was related to an inappropriate meaning of *spade*. The subjects were also required to state why the test word was inappropriate (e.g., stating something such as, "ACE refers to a playing card, but the type of spade suggested in the sentence is a garden tool"). Each of the 40 training sentences were presented four times during the training session.

Following the training session, all subjects performed the pictorial-suppression experimental task. As described before, the pictorial-suppression experimental task involves viewing an array of objects, and after viewing each array, verifying whether a test object was present. Trials which require suppression were those in which the test object is typically found in a scene that comprises the objects that were present in the array (e.g., a *farmer, pitchfork, pig, barn,* and *chicken* typically occur in a farm scene). By comparing the performance of the subjects who underwent the training with the performance of subjects who did not undergo the training, I was able to assess the transferability of the training intervention on the subsequent experimental task.

The data from this experiment suggest that training can transfer across domains. More specifically, preliminary data suggest that training to suppress interfering lexical information can improve performance in suppressing interfering pictorial information. Thus, being trained to suppress the interference caused by thinking of an inappropriate meaning of a homonym (e.g., the playing card meaning of *spade* as it is used in the sentence *He dug with the spade)* appears to improve the ability to suppress the interference caused by thinking of a typical but absent member of a pictorial scene (e.g., thinking that a *tractor* might be present in a farm scene).

Discovering that training in suppression in one domain leads to improved suppression in another domain has strong theoretical and practical implications. At the theoretical level, discovering that training in suppression in one domain leads to improved suppression in another domain supports the hypothesis of *one* general, cognitive mechanism of suppression. At the practical level, discovering that training in suppression in one domain leads to improved suppression in another domain demonstrates that persons' ability to suppress information can be improved, even without specific training in the domain in which suppression is required. Given the powerful role that suppression plays in numerous comprehension (and other cognitive tasks), this finding should lead to substantial applications for improving performance in various domains.

References

Biederman, I., Bickle, T. W., Teitelbaum, R. C., & Klatsky, G. J. (1988). Object search in nonscene displays. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 456-467.

Biederman, I., Glass, A. L., & Stacy, E. W., Jr. (1973). Searching for objects in real world scenes. *Journal of Experimental Psychology, 97,* 22-27.

Biederman, I., Mezzanotte, R. J., & Rabinowitz, J. C. (1982). Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive Psychology, 14,* 143-177.

Biederman, I., Teitelbaum, R. C., & Mezzanotte, R. J. (1983). Scene perception: A failure to find a benefit of prior expectancy or familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9,* 411-429.

Brownell, H. H., Potter, H. H., Bihrle, A. M., & Gardner, H. (1986). Inference deficits in right brain-damaged patients. *Brain and Language, 27,* 310-321.

Carreiras, M., & Gernsbacher, M. A. (1992). Comprehending conceptual anaphors in Spanish. *Language and Cognitive Processes, 7,* 281-299.

Carreiras, M., Gernsbacher, M. A., & Villa, V. (1995). The advantage of first mention in Spanish. *Psychonomic Bulletin & Review, 2,* 124-129.

Corbett, A. T., & Chang, F. R. (1983). Pronoun disambiguation: Accessing potential antecedents. *Memory & Cognition, 11,* 283-294.

Deaton, J. A., & Gernsbacher, M. A. (in press). Causal conjunctions and implicit causality cue mapping in sentence comprehension. *Journal of Memory and Language.*

Dell, G. S., McKoon, G., & Ratcliff, R. (1983). The activation of antecedent information during the processing of anaphoric reference in reading. *Journal of Verbal Learning and Verbal Behavior, 22,* 121-132.

den Heyer, K., Briand, K., & Dannenbring, G. L. (1983). Strategic factors in a lexical-decision task: Evidence for automatic and attention-driven processes. *Memory & Cognition, 11,* 374-381.

Faust, M. E., Balota, D. A., Duchek, J. A., Gernsbacher, M. A., & Smith, S. D. (1997). Inhibitory control during sentence processing in individuals with dementia of the Alzheimer type. *Brain and Language, 57,* 225-253.

Faust, M. E., & Gernsbacher, M. A. (1996). Cerebral mechanisms for suppression of inappropriate information during sentence comprehension. *Brain and Language, 53,* 234-259.

Foertsch, J., & Gernsbacher, M. A. (1994). In search of complete comprehension: Getting "minimalists" to work. *Discourse Processes, 18,* 271-296.

Garnham, A., Traxler, M. J., Oakhill, J. V., & Gernsbacher, M. A. (1996). The locus of implicit causality effects in comprehension. *Journal of Memory and Language, 35,* 517-543.

Gernsbacher, M. A. (1985). Surface information loss in comprehension. *Cognitive Psychology, 17,* 324-363.

Gernsbacher, M. A. (1989). Mechanisms that improve referential access. *Cognition, 32,* 99-156.

Gernsbacher, M. A. (1990). *Language comprehension as structure building.* Hillsdale, NJ: Erlbaum.

Gernsbacher, M. A. (1991a). Cognitive processes and mechanisms in language comprehension: The structure building framework. In G. H. Bower (Ed.), *The psychology of learning and motivation* (pp. 217-263). New York: Academic Press.

Gernsbacher, M. A. (1991b). Comprehending conceptual anaphors. *Language and Cognitive Processes, 6,* 81-105.

Gernsbacher, M. A. (1993). Less skilled readers have less efficient suppression mechanisms. *Psychological Science, 4,* 294-298.

Gernsbacher, M. A. (1994). Activating knowledge of fictional characters' emotional states. In C. A. Weaver, S. Mannes, & C. R. Fletcher (Eds.), *Discourse comprehension: Essays in honor of Walter Kintsch* (pp. 141-155). Hillsdale, NJ: Erlbaum.

Gernsbacher, M. A. (1995). The Structure Building Framework: What it is, what it might also be, and why. In B. K. Britton, & A. C. Graesser (Eds.), *Models of text understanding* (pp. 289-311). Hillsdale, NJ: Erlbaum.

Gernsbacher, M. A. (1996). Coherence cues mapping during comprehension. In J. Costermans, & M. Fayol (Eds.), *Processing interclausal relationships in the production and comprehension of text* (pp. 3-21). Hillsdale, NJ: Erlbaum.

Gernsbacher, M. A. (1997). Two decades of structure building. *Discourse Processes, 23,* 265-304.

Gernsbacher, M. A. (in press-a). Attenuating interference during comprehension: The role of suppression. In D. L. Medin (Ed.), *The psychology of learning and motivation.* San Diego, CA: Academic Press.

Gernsbacher, M. A. (in press-b). Group differences in suppression skill. *Aging, Neuropsychology, and Cognition.*

Gernsbacher, M. A., & Faust, M. (1991a). The mechanism of suppression: A component of general comprehension skill. *Journal of Experimental Psychology: Learning, Memory and Cognition, 17,* 245-262.

Gernsbacher, M. A., & Faust, M. (1991b). The role of suppression in sentence comprehension. In G. B. Simpson (Ed.), *Understanding word and sentence* (pp. 97-128). Amsterdam: North Holland.

Gernsbacher, M. A., & Faust, M. (1995). Skilled suppression. In F. N. Dempster, & C. N. Brainerd (Eds.), *Interference and inhibition in cognition* (pp. 295-327). San Diego, CA: Academic Press.

Gernsbacher, M. A., & Givón, T. (Eds.). (1995). *Coherence in spontaneous text.* Philadelphia, PA: John Benjamins.

Gernsbacher, M. A., Goldsmith, H. H., & Robertson, R. R. W. (1992). Do readers mentally represent characters' emotional states? *Cognition and Emotion, 6,* 89-111.

Gernsbacher, M. A., & Hargreaves, D. (1988). Accessing sentence participants: The advantage of first mention. *Journal of Memory and Language, 27,* 699-717.

Gernsbacher, M. A., & Hargreaves, D. (1992). The privilege of primacy: Experimental data and cognitive explanations. In D. L. Payne (Ed.), *Pragmatics of word order flexibility* (pp. 83-116). Philadelphia: John Benjamins.

Gernsbacher, M. A., Hargreaves, D., & Beeman, M. (1989). Building and accessing clausal representations: The advantage of first mention versus the advantage of clause recency. *Journal of Memory and Language, 28,* 735-755.

Gernsbacher, M. A., & Jescheniak, J. D. (1995). Cataphoric devices in spoken discourse. *Cognitive Psychology, 29,* 24-58.

Gernsbacher, M. A., Keysar, B., & Robertson, R. R. W. (1995, November). *The role of suppression in metaphor interpretation.* Paper presented at the 36th annual meeting of the Psychonomic Society, Los Angeles, CA.

Gernsbacher, M. A., & Robertson, R. R. W. (1992). Knowledge activation versus sentence
    mapping when representing fictional characters' emotional states. *Language and Cognitive
    Processes, 7,* 353-371.

Gernsbacher, M. A., & Robertson, R. R. W. (1994, November). *The costs and benefits of
    meaning.* Paper presented at the 35th meeting of the Psychonomic Society, St. Louis.

Gernsbacher, M. A., & Robertson, R. R. W. (1995). Reading skill and suppression revisited.
    Psychological Science, 6, 165-169.

Gernsbacher, M. A., & Robertson, R. R. W. (1997). *Parallel form affects sentence
    comprehension.* Manuscript submitted for publication.

Gernsbacher, M. A., & Robertson, R. R. W. (in press). The role of suppression in figurative
    language comprehension. *Journal of Pragmatics.*

Gernsbacher, M. A., & Shlesinger, M. (in press). The proposed role of suppression in
    simultaneous interpretation. *Interpreting.*

Gernsbacher, M. A., & Shroyer, S. (1989). The cataphoric use of the indefinite *this* in spoken
    narratives. *Memory & Cognition, 17,* 536-540.

Gernsbacher, M. A., & St. John, M. F. (in press). Modeling the mechanism of suppression in
    lexical access. In R. Klein, & P. McMullen (Eds.), *Converging methods for studying reading
    and dyslexia.* Cambridge, MA: MIT Press.

Gernsbacher, M. A., & Varner, K. R. (1988). *The multi-media comprehension battery* (Tech.
    Rep. No. 88-03). Eugene: University of Oregon, Institute of Cognitive and Decision
    Sciences.

Gernsbacher, M. A., Varner, K. R., & Faust, M. (1990). Investigating differences in general
    comprehension skill. *Journal of Experimental Psychology: Learning, Memory, and
    Cognition, 16,* 430-445.

Glucksberg, S., & Keysar, B. (1990). Understanding metaphorical comparisons: Beyond
    similarity. *Psychological Review, 97,* 3-18.

Haenggi, D., Gernsbacher, M. A., & Bolliger, C. M. (1993). Individual differences in situation-
    based inferencing during narrative text comprehension. In H. van Oostendorp, & R. A.
    Zwaan (Eds.), *Naturalistic text comprehension: Vol. LIII. Advances in discourse processing*
    (pp. 79-96). Norwood, NJ: Ablex.

Haenggi, D., Kintsch, W., & Gernsbacher, M. A. (1995). Spatial situation models and text
    comprehension. *Discourse Processes, 19,* 173-199.

Keele, S. W., & Neill, W. T. (1978). Mechanisms of attention. In E. C. Carterette, & M. P.
    Friedman (Eds.), *Handbook of perception* (pp. 3-47). New York: Academic Press.

Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words:
    Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology,
    90,* 227-234.

Oakhill, J., Garnham, A., Gernsbacher, M. A., & Cain, K. (1992). How natural are conceptual
    anaphors? *Language and Cognitive Processes, 7,* 257-280.

Palmer, S. E. (1975). The effects of contextual scenes on the perception of objects. *Memory &
    Cognition, 3,* 519-526.

Posner, M. I., & Snyder, C. R. R. (1975a). Attention and cognitive control. In R. L. Solso (Ed.),
    Information processing and cognition: The Loyola symposium (pp. 55-85). Hillsdale, NJ:
    Erlbaum.

Posner, M. I., & Snyder, C. R. R. (1975b). Facilitation and inhibition in the processing of signals.
    In P. M. A. Rabbitt, & S. Dornic (Eds.), *Attention and Performance V* (pp. 669-682). New
    York: Academic Press.

Prince, E. F. (1981). On inferencing of the indefinite-this NPs. In A. Joshi, B. Webber, & I. Sag (Eds.), *Elements of discourse understanding* (pp. 231-250). Cambridge, UK: Cambridge University Press.

Rayner, K., & Posnansky, C. (1978). Stages of processing in word identification. *Journal of Experimental Psychology: General, 107*, 64-80.

Rosinski, R. R., Golinkoff, R. M., & Kukish, K. S. (1975). Automatic semantic processing in a picture-word interference task. *Child Development, 46*, 247-253.

Simpson, G. B., & Burgess, C. (1985). Activation and selection processes in the recognition of ambiguous words. *Journal of Experimental Psychology: Human Perception and Performance, 11*, 28-39.

Smith, M. C., & McGee, L. E. (1980). Tracing the time course of picture-word processing. *Journal of Experimental Psychology: General, 109*, 373-392.

Swinney, D. A. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior, 18*, 645-659.

Tanenhaus, M. K., Leiman, J. M., & Seidenberg, M. S. (1979). Evidence for multiple stages in the processing of ambiguous words in syntactic contexts. *Journal of Verbal Learning and Verbal Behavior, 18*, 427-440.

Traxler, M. J., & Gernsbacher, M. A. (1995). Improving coherence in written communication. In M. A. Gernsbacher, & T. Givón (Eds.), *Coherence in spontaneous text* (pp. 215-238). Philadelphia: John Benjamins Publishing Co.

Tweedy, J. R., Lapinsky, R. H., & Schvaneveldt, R. W. (1977). Semantic-context effects on word recognition: Influence of varying the proportion of items presented in an appropriate context. *Memory & Cognition, 5*, 84-89.

van Orden, G. C. (1987). A rows is a rose: Spelling, sound, and reading. *Memory & Cognition, 15*, 181-198.

van Orden, G. C., Johnston, J. C., & Hale, B. L. (1988). Word identification in reading proceeds from spelling to sound to meaning. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 14*, 371-386.

Wald, B. (1983). Referents and topics within and across discourse units: Observations from current vernacular English. In S. Klen-Andrew (Ed.), *Discourse perspectives on syntax* (pp. 91-116). New York: Academic Press.

Wegner, D. M. (1989). *White bears and other unwanted thoughts*. New York: Penguin Books.

Wegner, D. M. (1992). You can't always think what you want: Problems in the suppression of unwanted thoughts. In M. Zanna (Ed.), *Advances in experimental social psychology (Vol. 25)* (pp. 193-226). San Diego, CA: Academic Press.

Wright, S., & Givón, T. (1987). The pragmatics of indefinite reference: Quantified text-based studies. *Studies in Language, 11*, 1-33.